6. Free Cooling via Waterside Economizer

Figure 1
Use of a Cooling Tower to Provide Waterside Economization

Data centers present an almost constant, 24-hour, internal cooling load. Free cooling can be provided via a waterside economizer, which uses the evaporative cooling capacity of a cooling tower to indirectly produce chilled water to cool the data center during mild outdoor conditions (particularly at night in hot climates). While a bin analysis using local weather data is required to properly assess the potential, free cooling is usually best suited for climates that have wetbulb temperatures lower than 55°F for 3,000 or more hours per year. It most effectively serves chilled water loops designed for 50°F and above chilled water, or lower temperature chilled water loops with significant surplus air handler capacity in normal operation. Often, existing data centers can capitalize on redundant air handler capacity with chilled water temperature reset controls to retrofit free cooling.

At least 3000 hours per year where wet bulb temperature is below:	Applicability
55°F	Free Cooling Specified Air Handlers or CRACs; Many Low Temperature Systems with Appropriate Chilled Water Reset
45°F	Retrofit Reusing Most Existing Air Handlers Designed for 42°F Chilled Water Without Need for Reset

Figure 2
Free Cooling Applicability

Principles

• While free cooling is operating, chilled water plant energy consumption costs are cut by up to 70%[1].

• Data centers require cooling 24 hours a day every day of the year – even when it is cold outside. This makes data centers very well suited to waterside economization. For example, in San Jose, free cooling would be expected to operate for over a third of the year, significantly reducing cooling bills and noticeably reducing chiller run hours and maintenance costs.

• Free cooling utilizing a waterside economizer can usually be economically retrofitted to existing chilled water cooled facilities.

• Isolation between the space air and outside air is not impacted by waterside free cooling, making it an alternative to airside economization when this is a concern.

• A flat plate heat exchanger is used to isolate the chilled water loop from the open tower condenser water to prevent fouling of coils.

• A low approach temperature cooling tower plant design is critical for best results. Use of redundant tower capacity can provide low approach temperature operation in a high-reliability 'spinning reserve' operation configuration.

• A traditional chiller is used to provide cooling during hot periods and as an always-available emergency backup. For a portion of the year, free cooling increases reliability by offering a non-compressor based backup to the traditional chiller, particularly at night when plant monitoring and staffing are liable to be lower.

Approach

Free cooling operates on the principle that during cool weather conditions, particularly at night, data center cooling loads can be served with chilled water produced by the cooling tower alone, entirely bypassing an energy intensive chiller. In this scenario, the cooling tower produces low temperature water. A heat exchanger is used to cool the building loop while keeping it clean and isolated from the relatively dirty cooling tower water. Free cooling reduces or eliminates chiller power consumption while efficiently maintaining strict temperature and humidity requirements. Other approaches to free cooling include: closed-circuit cooling towers (dry coolers) with glycol for climates with extended freezing conditions; earth coupled heat pumps; and/or independent free cooling loops serving dedicated free cooling coils in air handlers or Computer Room Air Conditioners (CRACs).

Data centers often have redundancy in their cooling tower plants. Through the use of VFDs and common condenser water header systems and/or sumps, the redundant tower capacity can be used to achieve a lower approach temperature. With variable speed fans, it is efficient to operate as many towers as the tower minimum flow requirements allow, maximizing the natural convective cooling while achieving a lower approach capability.

The majority of server computer equipment has a specified allowable operating range of 20% to 80% relative humidity, but ASHRAE recommends an operating range of 40 to 55%[2] and many data centers control the humidity more tightly. The minimum humidity is often controlled to prevent the risk of static discharge damage and the maximum humidity controlled to ensure a non-condensing environment and human comfort. Free cooling is well suited to such facilities, since it allows the space to be fully isolated from the exterior environment's humidity yet still reject heat without use of energy intensive compressor based cooling.

Free cooling can also offer an additional level of redundancy by providing a non-compressor cooling solution for portions of the year. In particular, free cooling can often provide a backup to compressor chillers during cooler nighttime hours when plant staffing may be lower. When the weather allows, free cooling replaces the complex mechanism of a chiller with a simple, non-mechanical heat exchanger.

Use of 'medium temperature' chilled water, in the range of 50°F and higher, maximizes the potential savings from a free cooling system. An efficient data center system is likely to be designed for use of chilled water in this temperature range already, since use of traditional (for office systems) 44°F chilled water is likely to result in uncontrolled dehumidification in high load areas (leading to either artificially low humidity or wasteful simultaneous dehumidification/humidification operation). A typical data center maintained at 72°F and 50% RH is susceptible to uncontrolled dehumidification when the supply water temperature is lower than 52°F.

New data center facilities that are specified for medium temperature chilled water can be designed. Medium temperature chilled water systems also reduce energy use and sometimes compressor sizing of normal chillers at the cost of somewhat larger coils and piping. In existing facilities, free cooling operation can be retrofitted and optimized through the use of an aggressive chilled water temperature reset. Typically, data centers have surplus air handler capacity, in the form of redundant units or overdesigned systems. The surplus capacity can allow the use of higher temperature chilled water than the original design through the use of a reset. A chilled water reset increases the chilled water supply setpoint in order to take advantage of the surplus capacity. The most efficient reset strategy is to have the central control system monitor the control valves on each air handler and/or CRAC and use the highest valve command (i.e., the most open valve) as the input to a control loop. The control loop resets the chilled water temperature, up or down, until the most open valve is at 90-95%. This approach automatically maximizes the use of installed coil capacity and automatically accommodates variations in data center load – increased cooling needs will result in the chilled water temperature being lowered if required to meet the load.

A cooling tower used for free cooling should be selected to provide an approach temperature (Leaving Tower Water Temperature minus Wet Bulb Temperature) between 5 and 7°F (see

Cooling Plant Optimization chapter). A lower approach temperature generally results in a physically larger tower since more surface area in the tower is required for the heat rejection process. A variable speed drive (VSD) for the fan motor should be used in a free cooling tower to minimize on-off cycling and maximize controllability and energy savings.

Figure 3 below shows a typical free cooling setup. Note that the free cooling heat exchanger is placed inline with the chillers. This allows for free cooling to support a portion of the load while the chillers can provide the last few degrees of cooling (referred to as integrated operation). An critical design element is that the chillers automatically and seamlessly provide 100% backup to the free cooling system. Failure of the free cooling system will simply result in the chillers seeing the full data center load; no control loop is required to switch between free cooling and chiller cooling, ensuring no loss in reliability from the addition of a free cooling system.

Figure 3
Free Cooling Loop Schematic

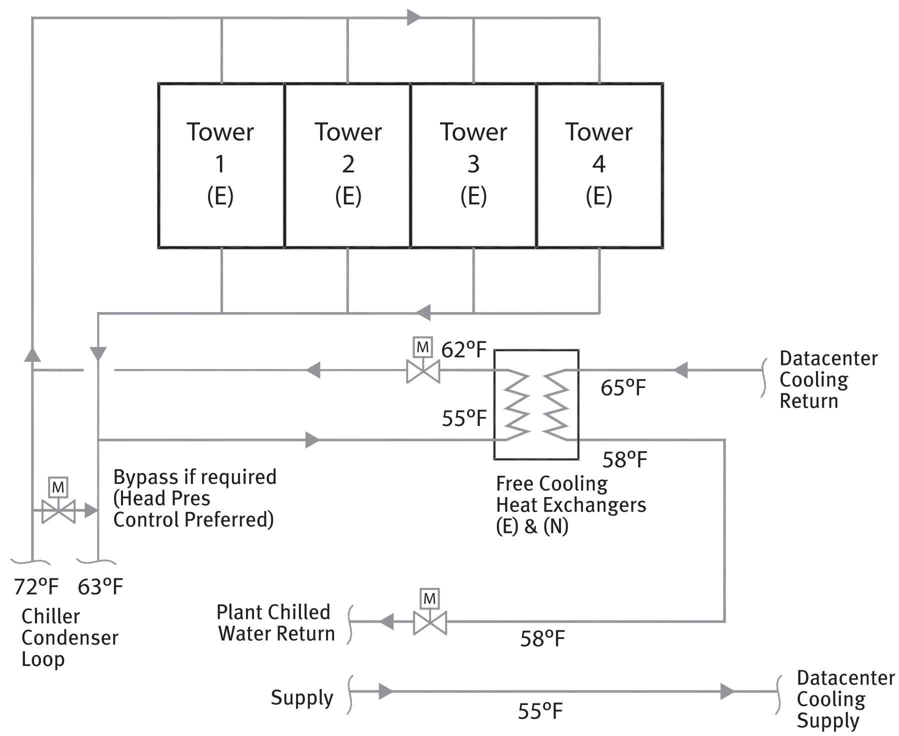

Free cooling requires that the cooling tower produce low temperature water, often lower than a chiller will accept for condenser water. There are two common design approaches to address this concern. One approach is to hydraulically isolate a tower and dedicate it to free cooling only. This is the best approach, but requires careful piping configuration and control valve placement and operation. A redundant backup tower can be provided with automatic isolation valves and used for free cooling. Since free cooling operates during low temperature weather conditions, the chilled water plant load is often low enough that even non-backup towers are available for free cooling use provided the proper automatic valving and a control sequence that gives the chillers priority for tower use (in case of free cooling failure) is implemented.

The other common approach is to share a single condenser water loop and towers between free cooling and the chillers by running the loop at a low, free cooling temperature and providing a bypass at the chillers, as shown in Figure 3. Locating the tower bypass (a standard feature in cooler climates) at the chiller end of the loop instead of at the cooling tower brings low temperature water into the main plant area, in many cases greatly reducing the cost of piping to implement free cooling. The bypass is used to mix warm condenser water leaving the chiller with low temperature condenser water to produce a suitable condenser water supply temperature – the standard tower bypass control loop. In cooler climates, a tower bypass is usually located directly next to the tower plant with an identical control algorithm to allow starting up the chillers during cold temperatures with a cold tower sump and chiller operation in very low temperatures. This approach is popular in retrofit situations or where the pipe run to the cooling towers is too long to economically allow a second set of pipes for free cooling. Some efficiency is lost by producing lower temperature water for the chillers than is used, but typically this is far outweighed by the reduced chiller compressor energy consumption.

Added costs for a waterside economizer result from controls, heat exchangers, and piping. Some installation will also incur additional costs for additional plant floor space or an additional pump. In typical critical facilities installation, no additional cooling tower capacity is required since the non-critical free cooling system components do not require any redundancy.

Benchmarking Findings/Case Studies

The psychrometric chart shown in Figure 4 represents a typical year's weather conditions in San Jose, California. Each hour is plotted as a point on the chart. A data center facility designed to utilize free cooling with 53°F water with a 10°F differentiated temperature could operate without any chillers for 2,800 hours per year. Partial cooling allows a total of 6,200 hours of cooling assistance for a predicted annual reduction in chiller energy use of 52%. A drier climate would yield larger savings.

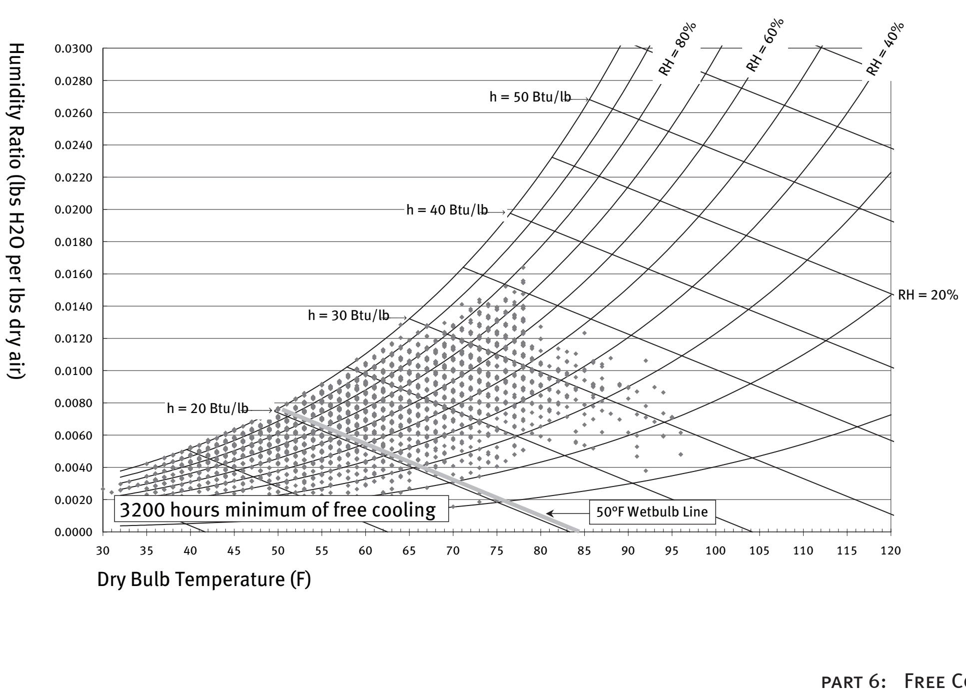

Figure 4
Psychrometric Chart for San Jose, California

Related Chapters

Centralized Air Handling   Air-side Economizer

Cooling Plant Optimization

References

1) Baseline study measurement of critical facility plants indicate that chiller power represents approximately 70% of total energy consumption.

2) ASHRAE, TC 9.9, Thermal Guidelines for Data Processing Enviroments

Pacific Gas and Electric Company®

Southern California Gas Company

A Sempra Energy utility℠

SDGE

A Sempra Energy utility℠

SOUTHERN CALIFORNIA EDISON®

An *EDISON INTERNATIONAL*® Company

Energy Design Resources provides information and design tools to architects, engineers, lighting designers, and building owners and develpers. Energy Design Resources is funded by California utility customers and administered by Pacific gas and Electric Company, San Diego Gas and Electric, Southern California Edison, and Southern California Gas Company, under the auspices of the California Public Utilities Commission. To learn more about Energy Design Resources, visit www.energydesignresources.com.

This Design Guidelines Sourcebook was developed by Rumsey Engineers. The following authors participated in the development of the guide: Larry Chu, John Weale, Peter Rumsey.